

Engineering a Successful Mission: Lessons from the Lunar Reconnaissance Orbiter

David F. Everett
NASA Goddard Space Flight Center
Code 599
Greenbelt, MD 20771
301-286-1596
David.F.Everett@nasa.gov

Abstract—Schedule pressure is common in the commercial world, where late delivery of a product means delayed income and loss of profit.¹² Research spacecraft developed by NASA, on the other hand, tend to be driven by the high cost of launch vehicles and the public scrutiny of failure—the primary driver is ensuring proper operation in space for a system that cannot be retrieved for repair. The Lunar Reconnaissance Orbiter (LRO) development faced both schedule pressure and high visibility. The team had to balance the strong push to meet a launch date against the need to ensure that this first mission for Exploration succeeded. This paper will provide an overview of the mission from concept through its first year of operation and explore some of the challenges the systems engineering team faced taking a mission from preliminary design review to pre-ship review in 3 years.

TABLE OF CONTENTS

1. INTRODUCTION.....	1
2. LRO OBJECTIVES	1
3. LRO RESULTS.....	1
4. PROGRAMMATIC ENVIRONMENT DURING DEVELOPMENT.....	3
5. CHALLENGES AND APPROACHES.....	6
6. I&T OVERVIEW	8
7. OBSERVATIONS AND LESSONS LEARNED	11
REFERENCES	12
BIOGRAPHY	12

1. INTRODUCTION

In January 2004, the President of the United States announced the “Vision for Space Exploration.” As part of this vision, the U.S. would send a series of robotic missions to the moon beginning no later than 2008. The first mission became known as the Lunar Reconnaissance Orbiter (LRO), and NASA released an Announcement of Opportunity (AO) for the LRO instruments in June of 2004, with a target launch of October 2008. Headquarters selected six instruments, each with a strong relationship to recently-flown payloads, in December 2004 and started funding for the mission in early 2005. Headquarters then added a synthetic aperture radar technology demonstration to LRO in April 2005. The Mini-RF development was significantly behind the other instruments, and its data rate and power

consumption were significant. The early vision of a relatively small spacecraft with a few instruments had turned into something larger, but it was still important to the new Exploration Systems Mission Directorate (ESMD) at Headquarters that LRO launch before the end of 2008.

2. LRO OBJECTIVES

LRO’s primary purpose was early reconnaissance in preparation for later return of humans to the moon. Certainly significant reconnaissance had been done in the equatorial regions of the moon prior to and during the Apollo missions, but long-duration stays on the moon would be difficult in the equatorial regions, due to the large swings in temperature (from -140 to +140 degrees C) and the two-week-long periods of darkness. The polar regions are much more hospitable, with a low sun angle, nearly continuous daylight on some mountain peaks, and permanently-shadowed craters that could harbor water. LRO specifically set out to characterize landing sites, identify resources, and characterize the radiation environment, particularly near the poles. The six instruments were specifically selected to accomplish these objectives, with robust overlapping of the measurements necessary to characterize the lighting, topography, and surface features of polar landing sites. LRO promised to improve our map resolution by several orders of magnitude.

3. LRO RESULTS

On September 16, 2010, LRO completed its mission for ESMD, meeting all mission objectives. The Science Mission Directorate has begun a science mission with LRO that will last another two years. The full details of the spacecraft and instruments can be found elsewhere [1]. Figure 1 shows the spacecraft in flight configuration, with the instruments labeled. What follows is a brief description of the LRO instruments and some interesting findings.

CRaTER

The Cosmic Ray Telescope for the Effect of Radiation (CRaTER) instrument is characterizing the radiation environment around the moon. The instrument includes tissue equivalent plastics which enable it to measure how much radiation would be deposited in human tissue.

¹ U.S. Government work not protected by U.S. copyright

² IEEEAC paper#1070, Version 1, Updated 2010:12:15

CRaTER has operated during a period when the galactic cosmic rays were at the highest levels ever recorded (due to low solar activity). It has detected unexpectedly high radiation levels near the surface of the moon, which are still under investigation.



Figure 1: LRO Spacecraft with Instruments Identified

Diviner

The Diviner Lunar Radiometer Experiment measures the infrared emissions from the lunar surface. Diviner has measured the coldest place ever recorded in the solar system. Some permanently shadowed craters near the lunar poles are below 35 K! Using Diviner's nine infrared channels, scientists have analyzed the global silicate mineralogy of the moon [2].

LAMP

The Lyman Alpha Mapping Project (LAMP) instrument measures the spectrum of the moon in reflected ultraviolet light. LAMP uses the galactic background Lyman alpha emissions to illuminate permanently shadowed regions. The instrument has discovered that permanently shadowed craters reflect less Lyman Alpha light than the rest of the lunar surface. The exact cause is still under investigation.

LEND

The Lunar Exploration Neutron Detector counts neutrons that come from the lunar surface. These neutrons are released when cosmic rays strike the material on the moon. Hydrogen is a very good absorber of these neutrons, so

LEND has identified regions of hydrogen concentration. The most surprising discovery about these hydrogen concentrations is that they are not always within permanently shadowed regions.

LOLA

The Lunar Orbiter Laser Altimeter is mapping the topography of the moon to an accuracy of better than 1 m with respect to the moon's center. LOLA uses a beam splitter and expander to create five separate 5 m laser spots on the ground. Taking five simultaneous measurements enables determination of not only the elevation, but also the slope in the region of the measurement. With well over 2 billion individual measurements, LOLA has created global topographic maps with horizontal resolution in the polar regions better than 50 m. This data set has allowed scientists to identify all craters greater than 20 km in diameter, then use that data set to measure the size distribution over time of objects that have struck the moon [3]. Figure 2 shows LOLA-measured topography in the vicinity of the Apollo 11 landing site.

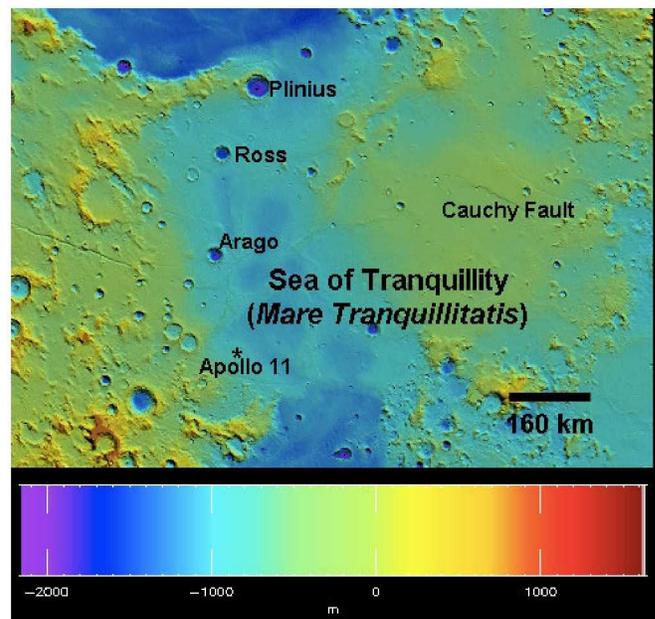


Figure 2: Image of LOLA Altitude Data (scale in m) [4]

LROC

The LRO Camera is composed of three separate cameras. The Wide-Angle Camera (WAC) provides 100 m pixel resolution in seven spectral bands (both optical and UV), with a 60 km swath width. The two Narrow-Angle Cameras (NAC) provide side-by-side coverage with 0.5 m pixels and a 2.5 km swath width, resulting in a total swath of 5 km. The cameras are line imagers, so they create pictures as LRO moves in its orbit. In addition to providing high-resolution imagery for possible landing sites, LROC has observed hardware left on the moon by previous missions, providing high-accuracy coordinates and context. An

LROC image of Lunokhod 1, a rover launched by the Soviets and last operating in 1971, has enabled Earth-based laser ranging off of a retro-reflector on the rover [5]. Previous attempts to find the reflector were unsuccessful, because the location estimate was off by several kilometers. Figure 3 shows the Apollo 17 landing site. Hardware and disturbed regolith (from the astronauts' activity) is visible.

Mini-RF

The Mini-Radio-Frequency is a synthetic-aperture radar with full polarization capability. This instrument was flown as a technology demonstration. It is proving its worth, generating images and maps of the moon's surface properties at S-band and X-band, with 15 m and 150 m resolution modes.

LCROSS Impact

The Lunar Crater Observation and Sensing Satellite (LCROSS) rode as a secondary payload on the LRO launch vehicle. LCROSS hung on to the upper stage of the Atlas V and directed it in to a permanently-shadowed crater near the

moon's south pole on October 9, 2009. LCROSS monitored the plume from the impact before striking the moon itself. Mission scientists used LRO data to identify the target, a permanently-shadowed region with a high hydrogen concentration. The LRO orbit was timed so that the spacecraft flew past the targeted site a few minutes after impact. The LAMP instrument identified volatiles from the moon in the plume, and the Diviner instrument imaged the heat left over from the impact on subsequent orbits. Analysis of the LCROSS data indicates the presence of water in the ejected plume.

4. PROGRAMMATIC ENVIRONMENT DURING DEVELOPMENT

Systems engineering is more than simply a technical process. Certainly technical requirements and constraints drive a system design, but programmatic considerations, such as cost, schedule, and political perceptions, will also exert significant influence on the technical decisions the team makes. A brief description of this environment will make the systems engineering approach much more understandable. The in-house development team at

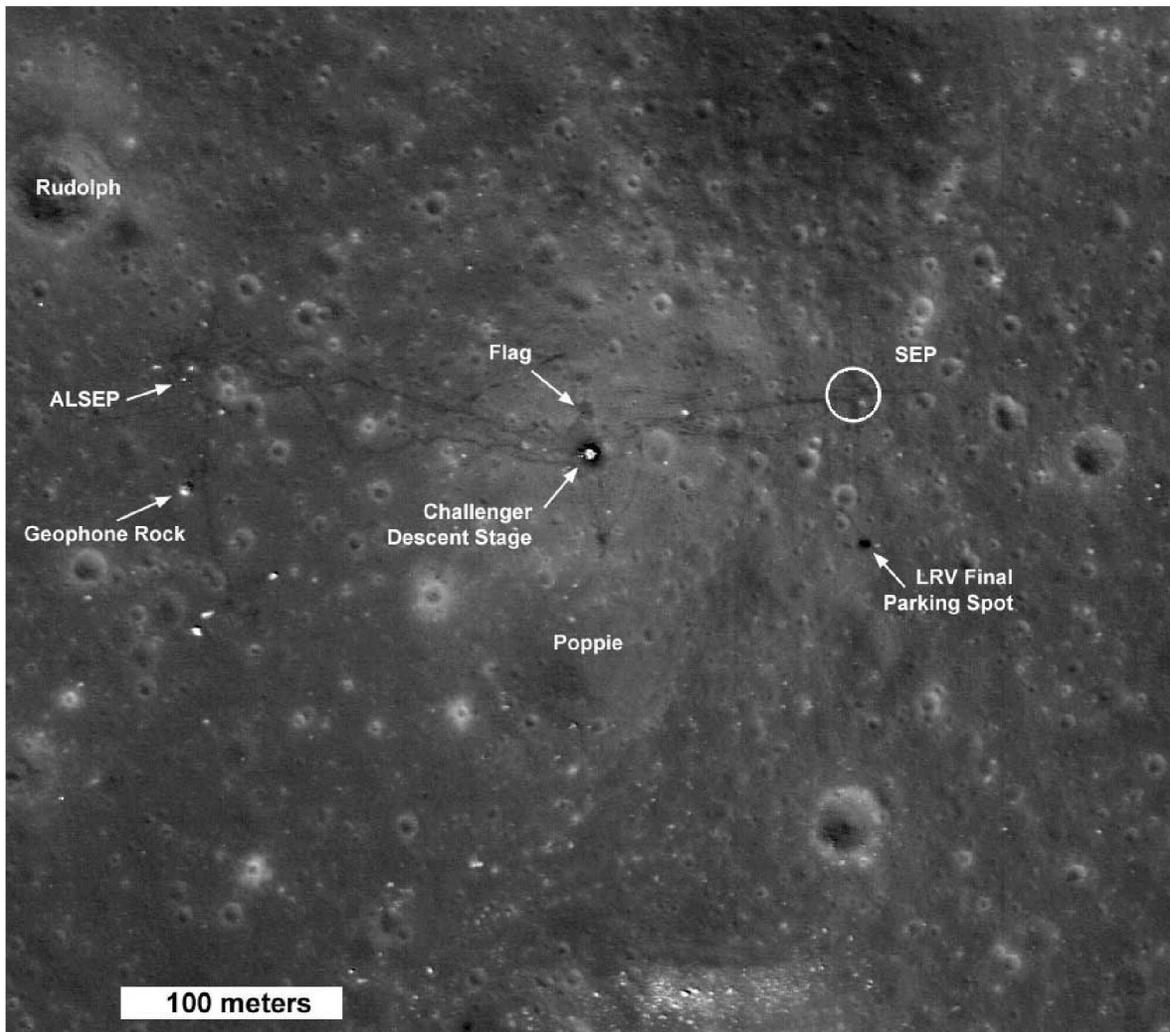


Figure 3: LRO Camera Image of Apollo 17 Site [6]

Goddard Space Flight Center (GSFC) developed the conceptual design of LRO, prior to the instrument AO release. Because the instruments were proposed to this concept, it led to the primary mission design constraints. The only major change to these initial constraints was a change in launch vehicle shortly before the mission Preliminary Design Review (PDR). The biggest programmatic driver of the mission was the schedule constraint, although several other factors came in to play.

Development Organization

NASA decided at the beginning to designate LRO as an in-house mission at GSFC. GSFC generally takes on a certain amount of in-house work, where NASA engineers lead the design and build of the flight hardware—designing circuit cards, software, and structures, with the spacecraft assembled and tested in GSFC facilities—in order to maintain technical expertise within the organization. LRO provided significant training opportunity for a number of GSFC engineers, and NASA gained the added benefit of not locking into a large spacecraft contract during the rapidly changing political environment of the Exploration Initiative startup (see *Changing Program Office* below). The in-house leadership team included a number of individuals who had worked on the Small Explorer missions in the 1990's. This experience with rapid, low-cost, highly-reliable, single-string spacecraft provided an excellent foundation for the LRO development.

At its peak, the LRO core systems team included 15 GSFC employees and on-site support contractors. These engineers looked across the project from different perspectives, including specialties such as modes & operations, launch vehicle interface, avionics, manufacturing, verification, and requirements management support. The extended systems team included the core team and all of the subsystem leads. Although the subsystem leads primarily focused down into their subsystems, they were also expected to understand how their piece of the system interacted at the system level. The subsystem leads participated in weekly project meetings and major architecture decisions, and they were key contributors to the risk management process. The extended systems team provided the technical leadership for the entire LRO team.

Although the systems team had a hierarchical structure, with the Mission Systems Engineer (MSE) at the top, team members were encouraged to make decisions within their area of responsibility, only elevating decisions when the impact crossed multiple boundaries. For simple interfaces or interactions, subsystem leads worked the details among themselves, with a member of the core system team approving the solution. For major architecture decisions, such as the one described in the *Launch Vehicle Change* section below, the extended systems team provided the technical, schedule, and cost perspectives that both the MSE and the Project Manager (PM) needed to make the final decision. The PM strongly encouraged diverse perspectives,

and selected team members who complemented each other. The MSE worked diligently to ensure team members throughout the organization spoke up when something seemed amiss, following through to understand the situation and provide feedback to the team on the resolution.

Design Constraints

Early on, the LRO team chose a 50 km, polar operational orbit. This orbit enabled very high-resolution imagery and laser altimetry across the entire moon. For orbits lower than this, the orbit maintenance must be done more often than monthly, due to the asymmetry of the moon's gravitational field. This orbit became a constraint on the flight system. From 50 km altitude, the moon fills almost half of the field of view, so the transition from noon to midnight in a 113-minute orbit really drove some aspects of the thermal design. The thermal considerations were probably the biggest change for the LRO instruments from their heritage, mostly Mars, designs. The polar orbit is inertially fixed, so it experiences all sun angles over the course of a year, resulting in a two-axis gimbal on the solar array.

Since the orbit maintenance at 50 km is fairly high, the instrument AO defined a 14-month mission, with the first two months used for instrument commissioning. A year of operations provided enough time to get thorough, high-resolution coverage at the poles. Since LRO was intended to be the first in a series of missions, and since it would be a short mission, LRO was designated as reliability class "C", meaning that it would be single-string. The initial launch vehicle was the 3-stage Delta II. The third stage is spin stabilized, and the total capability to a lunar trajectory was 1480 kg. The change in launch vehicle is described later.

Changing Program Office

The LRO team formulated the mission and spacecraft design during a turbulent period. NASA's Exploration Systems Mission Directorate (ESMD) at Headquarters was just getting started, and the Robotic Lunar Exploration Program moved from Goddard Space Flight Center to Ames Research Center and then to Marshall Space Flight Center. LRO's original design approach was "design to cost", but after selection of the six instruments and technology demonstration payload, the LRO team's direction was "design to requirements". Cost became less of a concern to ESMD for a while, and there was a strong desire to fly Mini-RF. As the LRO approached confirmation review, and the full cost of the increased requirements became apparent, the team was again directed to focus on cost. This pinch on cost around the time of the PDR helped the team focus on what expenses were critical and which were not. It turned out to be an effective measure to enable the LRO management to pull together a focused project plan, but with sufficient funds to deliver on schedule.

Schedule

Because this was the first mission for ESMD, the organization’s political credibility became linked to LRO’s successful launch in the Presidentially-mandated year of 2008. There was enormous pressure on LRO’s project management to hold schedule. After the Critical Design Review (CDR), ESMD offered the project additional funding if it could save schedule. Because all of the major contracts were already in place, very little additional funding could be spent effectively, but the project was able to apply some money to some key areas, like an interface test for the transponder. This spending had the effect of keeping all components on schedule, preventing a single delayed component from holding back the entire development.

The Independent Review Team expressed concern with the schedule at each of the reviews. The LRO engineers frequently expressed concerns with meeting schedule. The biggest challenge for the project management and systems engineering team was the balance of schedule against technical risk. Early on, someone on the team found in a fortune cookie the note in Figure 4, which summed up the feelings of many on the engineering team. The challenge became finding ways to get things done faster without compromising the technical integrity of the mission. Before getting into the details of how we did this, a description of the launch vehicle change will highlight several aspects of the LRO systems approach, demonstrating the large influence schedule pressure put on the system architecture. The actual schedule executed by LRO is shown in Table 1.

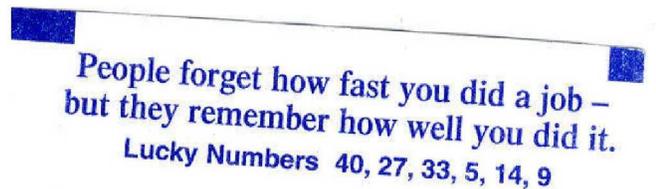


Figure 4: LRO Fortune Cookie

Launch Vehicle Change

In September 2005, the propulsion team identified a risk associated with the nutation time constant (NTC) of LRO’s propellant tank. In order to support the mission design, including the lunar orbit insertion burn, the delta V for the mission was 1258 m/s. Because of the short development schedule, the team baselined a mono-propellant hydrazine system, resulting in nearly half of the launch mass being fuel. The tight mass and volume constraints of the Delta II led to a custom tank design for the mission. With the spin-stabilized upper stage, the Delta II had strict requirements on the system’s NTC, so that sloshing of the fuel would not cause nutation beyond the capability of the Delta’s control system. LRO’s combination of mass and fuel fraction was outside the experience base of the Delta launch vehicle—no

one had ever flown as massive a spacecraft that was almost half fuel on this rocket.

Table 1: LRO Actual Milestones

AO Released	Jun 2004
Funding start	Jan 2005
System Requirements Review	Aug 2005
Launch vehicle changed to EELV	Nov 2005
Preliminary Design Review	Feb 2006
Confirmation Review	May 2006
Critical Design Review	Nov 2006
Start of spacecraft integration	Jan 2008
Pre-Environmental Review	Jun 2008
National launch priorities slip LRO to 2009 launch	Jul 2008
Thermal Vacuum start	Oct 2008
Environmental testing complete	Jan 2009
Pre-Ship Review	Feb 2009
Launch on June 18, 2009	Jun 2009

The team expected that they could eventually come up with a tank design that would meet the NTC requirement, but we would not know for sure until the spring of 2006, after some drop-tower testing. For LRO, waiting six months to determine whether or not to change tank configurations was unacceptable. The systems team immediately investigated several options including a bi-propellant system, a different launch vehicle without a spinning upper stage, and a solid rocket motor for the lunar orbit insertion burn. Because a change in launch vehicles was outside the decision space of the project, and because the bi-propellant system would take too long to develop and it may still have a NTC issue, we selected the solid rocket motor option and began designing the stage and redesigning the spacecraft structure. We then briefed ESMD in November 2005.

At the ESMD briefing, we pointed out that a change to the non-spinning Evolved Expendable Launch Vehicle (EELV), in addition to solving the NTC issue, would add sufficient margin to the mission to use surplus propellant tanks from the X-38 program, and it would add additional capability to launch a secondary payload. Right at the briefing, the Associate Administrator decided that the advantages of lower development risk for LRO (existing tanks) and capability for an extra payload (later became LCROSS) easily offset the higher cost of the larger launch vehicle. He switched LRO to the EELV and directed the team to optimize the design for launch on this vehicle.

The new vehicle with the existing tanks meant that the mechanical team needed to start over with their design. We had learned quite a bit about the thermal challenges of this orbit during the preliminary design work that had already been done. So the systems engineering team provided some guiding principles for the new structure design:

- (1) We set the mass limit for the system using the fuel capability of two X-38 tanks. With our delta V, the extra dry mass capability for LRO was about 174 kg, but 114 kg of that was consumed with the heavier tanks and the accommodations, leaving a dry mass-margin increase of 60 kg. Our overall margin at the time went from 24% to 28%.
- (2) We created a modular spacecraft configuration that would enable parallel assembly, making up the schedule lost by changing configurations. All of the avionics were mounted to a single deck for table-top integration. The propulsion system could be built up as a separate assembly, with all of the plumbing and thrusters integral to that assembly. We had a graphite-composite instrument module for accurate payload pointing stability. During system-level Integration and Test (I&T), we had three teams working simultaneously. The propulsion module was pressure proof-tested by itself, while we integrated avionics in a clean room and installed heaters and thermostats on the instrument module in another location.
- (3) We coupled significant mass together in order to minimize the thermal transients in the noon-midnight orbit. The avionics deck contained heat pipes to spread the heat and transfer it to a zenith-facing radiator, and the wheels were coupled into the same radiator. The instruments were thermally isolated with a clear view of space away from the sun. Coupling the mass together also reduced heater power, especially during the long lunar eclipses, when the moon passes through the Earth's shadow.

The structural re-design put the mechanical and thermal teams significantly behind when we held our mission Preliminary Design Review (PDR) in February of 2006. The propulsion team had originally planned to procure an integrated propulsion system, but the use of in-hand components shifted this effort in house. This change greatly increased the workload for the propulsion organization, but it also provided some good opportunities for hands-on experience. The other subsystems were ahead of PDR level at this time, so we were able to hold a successful review, with mechanical, thermal, propulsion, and attitude control performance peer reviews in May 2006 covering the necessary system-level issues. By CDR in November of 2006, these subsystems had caught up with the rest of the system. Changing structures was painful for a while, but in the end it proved to be quite a blessing, with the modular integration and the coupled thermal design.

As a side note, the use of the existing propulsion hardware was not as inexpensive as one might expect. The LRO team needed to review all of the requirements and paperwork associated with the X-38 hardware. The attitude control thrusters were not specified to handle the throughput that LRO required, so the team needed to procure additional testing for those components. All of the details necessary to

ensure application of existing hardware to a new mission add up to a significant effort, but this effort is essential to ensure proper operation in the system.

5. CHALLENGES AND APPROACHES

Designing for parallel development with a modular design was one way that the LRO team was able to get things done faster without compromising technical integrity. Other approaches included working harder, adding more people, testing early, managing risk, challenging processes, focusing on people, and making decisions effectively.

Work Harder

Obviously working harder can get things done faster. The challenge is motivating people. The LRO management was very good at this. In Daniel Pink's book [8], he describes three components to motivation: autonomy, mastery, and purpose.

Autonomy—The project management and the systems team invited team members at all levels of LRO to promote their best ideas in support of the mission as a whole. The team members felt ownership for these solutions, and felt they had control over their own destiny. We delegated decisions to lower levels as much as possible, with systems engineering holding the big-picture view, facilitating communication, documenting design decisions, and verifying system-level performance. Team members understood their control over mission success and were, therefore, motivated to excel.

Mastery—We challenged our experts in a constructive way. The systems engineers used basic physics to double check solutions, and we expected the experts to be able to explain their results in those terms. We encouraged discipline leads to represent their own perspective, but with an eye to the system impact. We expected ownership and probing inquiry of anomalies and design issues.

Purpose—In some respects, this was the easiest motivator. We were headed to the moon, paving the way for new exploration. It was an engineer's dream. But the moon doesn't excite everyone. The technician putting a part on a board does not always see the bigger picture, so we frequently reminded everyone of the criticality of every piece of the system. We walked around and talked with the technicians, explaining the excitement of the mission and the importance of each piece. And when we completed a review or met a milestone, we held a celebration. We created a sense of team that became like a family, where we all enjoyed each other's successes, because we each knew that each victory got us closer to mission success.

Add More People

In some areas, additional people helped us move faster. The mechanical team was able to bring in a number of additional

designers to work the details of piece parts and a few additional leaders to take charge of the modular structure components. This is primarily how they were able to make up for the time lost in the launch vehicle change. But if there are not enough leaders, extra manpower will result in idle hands, potentially distracting other team members. Be clear about the responsibilities of new team members.

Test Early

We conducted interface tests with each instrument and with most other components as soon as breadboards were available. At nearly every one of these tests, we discovered problems that would have been a disaster during system-level integration. For example, we added an interface test for the transponder, and that test discovered two signals were reversed, a reset had not been implemented per the specification, a software error on the spacecraft side caused repeated strobing of the hardware reset, the auxiliary command port worked better with a square wave rather than the sine wave specified, and the telemetry database was incorrect. These problems were trivial to fix on the breadboard system (most were corrected during the test) and they cost nothing to fix on the flight system, since it had not yet been built. Had we not done the interface testing, we would likely have had some or all of these problems show up during the flight integration in February 2008. It would have taken a couple of months to correct the issues on the flight unit, and our entire spacecraft integration would have been set back. The time saved with this and every other interface test was well worth the time required to execute the tests. During all of LRO's system integration, we found only one interface problem (a reversed 1553 transformer), an incredible testament to the benefits of interface testing.

Manage Risk

The LRO risk management process worked very well. The mission systems engineer and the risk manager met with each subsystem lead individually every month. During this meeting, we would update the subsystem's risks and we would discuss anything that might be worrying the lead. The relaxed atmosphere, with the meeting conducted at the lead's office, led to frank discussions of issues and potential problems. By meeting monthly, we were able to keep the database current and satisfy monthly reporting out from the project, but more importantly, we were able to identify and begin to mitigate things that would have become serious problems without project-level action.

Our monthly risk management board meetings at the project level provided a forum for discussing the risks we identified. These discussions gave the project managers critical insight, and it also provided a forum to discuss at the project level where we should best spend our resources. Because of our schedule, if someone identified a critical risk, we would not wait for a monthly meeting to address it. We would gather the necessary engineers and managers and begin risk mitigation immediately. The previously

discussed risk associated with the propellant tank NTC and its impact on the launch vehicle is a good example of LRO's risk management in operation for a critical risk.

A more routine example of LRO risk management involves the late delivery of a part on the Command and Data Handling (C&DH) power supply board. At a monthly meeting, the lead identified that a part might arrive late, delaying the integration of his flight hardware. This was an issue the lead had raised with the project-level parts support, but systems engineering was not aware that there might be a delay until the monthly meeting. The project decided to build up a flight spare card using a slightly lower-quality part for this one area. This spare card enabled the C&DH lead to continue with his integration, while our quality engineer started investigating the risk associated with flying this lower-quality part. All of this action bought enough time so that when the flight part did actually get delayed, we could keep moving while we waited for its delivery. In the end, the flight part came in early enough to integrate with the flight system before qualification testing. Had our lead not mentioned his concern during a risk meeting, we would have ended up delaying the entire flight integration by a month or more.

Challenge the Standard Processes

The LRO team habitually questioned why we did things the way we did them. If we did not have time for a particular test, we did not simply skip the test, we figured out why the test was required and then figured out a better way to accomplish the same thing. Procurements are an area that generally has well-established processes, some of which are based on previous procurements, whether or not they were successful. The LRO systems team did not leave procurements to chance. We had significant systems support for each of our major purchases of spacecraft hardware, ensuring the proper environmental requirements and good interfaces. We had a manufacturing engineer on our team who worked with each vendor after contract award. This engineer tailored our Statement of Work and our specification in order to use the vendor's processes as much as possible without sacrificing quality. This engineer checked the preliminary parts lists for any potential risks, and he ensured that the vendor planned to do the right tests. These early meetings made sure that the vendor understood the LRO requirements and understood the importance of the component to mission success. Because of this collaborative approach, we saw almost no problems with procured components at the time of delivery. The pre-ship reviews for each component went remarkably smoothly.

Focus on the People

We solved problems quickly by bringing our full talent to bear. LRO systems engineering actively encouraged diverse perspectives and minority opinions. Although these are popularly stated goals, it is not easy to get this input from most people. Especially when there is a close sense of team,

individuals are reluctant to rock the boat or express an unpopular position. Frequently people will not express an opinion because they do not want to be a bother or they do not realize the full, system-level impact of their concern. And sometimes people do not think their voice will make a difference. The LRO systems team countered these effects by actively listening to concerns, and following up with at least an investigation. It is important to follow up with the individual to describe the disposition of the concern, even if the best course of action is to proceed as planned. If an individual hears why his concern was set aside (perhaps other concerns were more pressing), then he is more likely to raise a concern in the future, as opposed to the situation where it appears to him that his concern was totally ignored.

It is not easy to encourage everyone to let you know what is bothering them. Some people will bring up too many issues (the constant worriers); some people's styles do not match; some people have less experience, so they do not know what might be a concern at the system level; and some people just think differently. But it is the different perspective that adds the strength to the team. Once launched, a spacecraft is usually irretrievable. The systems team must find all of the fatal flaws before that launch. Different perspectives fill in the blind spots, helping us to see where things might go wrong.

A great example is the LRO coarse sun sensor (CSS) circuit design. During system-level testing of the CSS's, the engineers who were performing the test noticed a very small discrepancy in the current flow (on the order of a few micro amps). The mission systems engineer encouraged these engineers to keep digging into the source of the discrepancy, even though he expected it was just some bookkeeping mistake. The engineers scheduled multiple tests on the spacecraft, to the point where the I&T manager started to complain about their chasing of what appeared to be nothing, but the systems engineer defended them. Eventually, they tracked the issue to a flaw in the circuit that read the CSS's. Although subtle in the test, this flaw was big enough that the sensors would have been essentially useless in flight. We needed to remove the C&DH box and change some resistors. Had these engineers dropped the issue, LRO would have launched with non-functioning sensors that were critical to acquisition of the sun. We might have lost the mission.

Make Decisions and Move On

There is no fool-proof method to make multi-parameter decisions with high-stakes risks. LRO could have quickly fallen behind if we took too long to make these decisions, and we could have failed spectacularly if we got an important one wrong. The LRO team, and particularly the systems engineering team, handled important trades by taking input from multiple people (getting those diverse perspectives), analyzing what could be analyzed, and then picking a path using *engineering judgment*. The human mind is very good at taking diverse inputs and synthesizing

them into an understanding of the situation—we do it all the time with the diverse input of our five senses. If we try to bypass this capability with some rote process that looks at each parameter individually, we are more likely to end up with the wrong answer. The key is getting perspectives from multiple people to reduce the bias. With that input and a list of the factors that play into the choice, the decision maker has what he needs.

The LRO team did not spend a lot of time looking for other options if we found one that met schedule, cost, and performance requirements. By choosing the first viable solution, the team did not always reach an “optimal” decision, but since the schedule was tight, it was usually quicker to take that path. The biggest challenge with this approach was presenting the results to review teams. Reviewers expect to see evidence of careful, exhaustive trade studies, with clear, analytic rationale. We did perform analysis when it made sense (such as determining the mass margin for each of the propulsion options), but we did not have the luxury of spending six months analyzing the nutation time constant of our tank. The review teams were generally satisfied with the list of criteria we used to make the decision—they could then use their own judgment to assess the reasonableness of the decision.

Decisions cannot wait for complete information. One of LRO's major trades involved the data storage system. The original baseline system used commercial hard drives. Radiation testing demonstrated that several parts in the drives might have issues in the space environment. The experts believed that we may be able to come up with some circuitry to handle the issues. We needed to decide whether to continue down this path, or create a new RAM-based Data Storage Board (DSB) design. We only had a notional design for the DSB solution, with very little engineering to back it up. But our engineers had built similar boards before. The mission was beyond PDR, and the rest of the C&DH was well on its way to the CDR. LRO could not wait for a preliminary design of the DSB's before choosing. With the information available, we decided to make the switch to the DSB, and we did not look back. As with all LRO decisions, we proceeded at full speed on the chosen path, only considering new options if we hit an obstacle. The DSB board design caught up with the rest of the C&DH, and the software team did fine with the change in mass storage (although there were a few bumps there, too).

6. IMPLEMENTATION

Component Build

The standard systems engineering process generally does not say much about the role of systems engineering during component build. The requirements have been allocated, and the systems engineering team is planning for system-level verification and validation. But the reality demands that the systems engineer stay connected with the team, looking for mistakes and misunderstandings. During this

phase of the LRO development, nearly every component had some sort of *unforeseeable* problem, with technical risk-vs.-schedule or subsystem-vs.-subsystem decisions required. At the project level, LRO followed over 50 issues between CDR and the start of I&T. Requirements were misunderstood, performance failed to meet expectation, parts failed, drawings were wrong, and components were misassembled. Sometimes we needed to bring extra engineering talent to bear on a problem, sometimes we needed to change a requirement in one area to pick up the slack of missed performance in another area, sometimes we simply waived a requirement because we had enough margin elsewhere, and sometimes the component team worked through their own problem within their own margin. The systems team was very active in this period, ensuring that all components would work together when the system was assembled. The biggest advice from this period was an often-spoken but never-written rule on LRO: wait 24 hours before reacting to bad news. This pause allowed the component team to properly assess the situation and it kept the entire team (and upper management) calm. Nearly always, things looked better the next day. The component team understood the situation, and they had time to generate solution options.

System-Level Integration and Test

LRO's avionics table-top integration (see Figure 5) went very smoothly. Having easy access to connectors made it very easy for the technicians to walk through the safe-to-mate procedures. The propulsion module (see Figure 6) integration started almost a year before the rest of the system integration. Technicians performed over 220 welds and installed as many thermostats for heaters. Parallel assembly of the propulsion module, avionics module, and instrument module converged in March 2008 (see Figure 7). The LRO team worked the rest of I&T on scaffolding.



Figure 5: LRO Avionics Table-Top Integration

A major piece of the strategy for building a single-string spacecraft on a short schedule was the test program. The LRO team planned a robust test program in order to ensure we caught any serious flaws before launch. The single-

string transponder was used for nearly all commanding and telemetry during I&T, most of the time through the RF link operating near threshold, so any changes in gain would be obvious. We performed the standard EMI, vibration, shock, and thermal vacuum tests. We conducted numerous operations tests and flight simulations, flowing data from the instruments all the way through the science data processing. By the time it launched, the LRO flight system had logged over 3400 hours of powered operation.



Figure 6: LRO Propulsion Module



Figure 7: LRO Y-Panel Installation

Launch Date Slip

In July 2008, with LRO already starting environmental tests, other launch priorities within the government led to a delay in LRO's target launch date. We initially targeted the end of February, but additional issues with the launch vehicle and payloads ahead of LRO slipped the launch out to June 18, 2009. The team used the extra time to increase the amount of system testing. We spent quite a bit of the time with operations testing, ensuring that we had practiced for major contingencies. Although we were still on track to make a 2008 launch at the time of the slip, we would likely

have missed the target by a month or two because of a mistake in the thermal design of the solar array gimbal (more details later).

Closeout of Paperwork

The systems team and the mission assurance engineers worked together to ensure clear and complete documentation. We established acceptance reports for components. These reports provided pointers to all of the necessary documentation, such as verification reports and performance measurements. They were signed by the component owner, the subsystem lead, and the cross-cutting systems engineers (mechanical systems, avionics systems, and thermal systems). These acceptance reports ensured that all necessary documentation was in place and everyone had reviewed their piece of it.

We did not consider a requirement verified until it had a closed work order, documenting a test or inspection, or until it had a report in the configuration management system, documenting the analysis. Work orders were not closed until all associated problem reports were closed and all paperwork was complete and in the database. LRO was moving very fast. Our attention to paperwork detail helped ensure that we did not let a critical item slip through the cracks.

Launch Site Operations

The biggest challenge with launch site operations is the attention. Senior managers suddenly get very nervous, and bad news spreads like wildfire, even if it is not true. Any mistake at the launch site is under the spotlight, so it is a time for extra care. This extra care frequently results in engineers spotting subtle quirks that have been there all along. The systems team needs to remain calm and work diligently through each of these issues. LRO found several quirks during this period.

At the launch site, the project loses control of the schedule. For a team that had been as focused on schedule as LRO, this fact was particularly annoying. Launch vehicle and range issues are generally not under the control of a project, so the spacecraft team must stay flexible, ready to juggle the work and move when necessary.

Launch operations is a time of camaraderie and stress. Part of the team is stationed with the spacecraft, working the flight system issues, while the rest of the team is at home, preparing for operations. This separation does create some division within the team, so LRO management put some extra effort into connecting the team. A coordination meeting for the operations team at the launch site, combined with a tour of the launch facilities, helped everyone better communicate and prepare for launch day. Figure 8 shows the LRO launch on June 18, 2009.

Post-Launch Operations

LRO used a direct-insertion trajectory, following a 5-day course to the moon. Because of very low errors in the Atlas V launch insertion, the mid-course correction 24 hours after launch was only 1.3 m/s (we had budgeted 25 m/s). The 40-minute lunar orbit insertion burn on June 23 changed the spacecraft speed by 555 m/s. The pre-launch rehearsals paid off, with an efficient team guiding the system through an almost flawless check-out and lunar capture. We did encounter a couple of safe mode transitions early in the mission—one due to overly tight safing limits and the other due to unexpected behavior of the star trackers on the transition from lunar occultation. The simulations never completely match the real environment. There is much to learn in the first few weeks of the mission, so our 24/7 staffing was essential. Later in the mission, operator errors also precipitated a couple of safe mode transitions (sequence error on an off-nadir slew, and a sequence error on a data dump just before a station-keeping maneuver; we have since implemented more rigorous review of non-routine sequences); in every case, the spacecraft took care of itself as designed. The instrument commissioning took a couple of weeks longer than expected; we had consciously put less attention on this phase pre-launch in order to ensure full readiness in other areas. Overall, LRO operations has been very smooth. Between mid-September 2009 and mid-September 2010, the spacecraft operated in nominal science mode over 96% of the time. The operations center has collected a total of over 39,000 Gbytes of raw data in over 500,000 files.



Figure 8: LRO Launch on the Atlas V Vehicle 20

7. OBSERVATIONS AND LESSONS LEARNED

Looking back at the LRO development, the team did some very important things very well, and we, of course, made some mistakes.

Mistakes

On any complicated development, with many different people involved, expect mistakes! We all make mistakes; the challenge is to avoid the fatal ones. Treat mistakes and failures as a normal part of the job. Keep things out in the open, clean up, and move on. It is critical to understand the cause of a mistake, but avoid blame. It does not help to punish someone who is trying his best to succeed. You can be sure that a conscientious engineer will be much more diligent after making a mistake that has caused much grief.

The LRO systems engineering management underestimated the system-level complexity of the High-Gain Antenna System and the Solar Array System. These systems both involved deployable systems with two-axis gimbals. The interplay of the electrical, mechanical, RF, and thermal really needed a full-time systems engineer, but we did not realize that until late in the build phase. On our solar array gimbal, we almost missed a critical test—running power through the harnessing during thermal vacuum testing. Had we not picked up this test at the system level (thanks to a diligent reviewer), an incorrect analysis of harness power dissipation would have led to overheating of the system in flight—we would probably have had a crippled gimbal.

Another error of the systems engineering management was an underestimate of the warning signs coming from some of our overloaded engineers. High performers sometimes do not take a break when they should, and they sometimes take on more than they can handle. It is the leader's job to spot these situations and take action. But if it is one of the leaders getting overloaded, sometimes we don't take the action we should. We saw some of the warning signs, and we did generally make sure the engineers closest to the hardware got enough down time. But we used very little of the extra time from the launch delay to rest our leaders. In hindsight, I think we ended up being less efficient in some areas, with some overloaded engineers becoming bottlenecks, but we were very careful to avoid impacting the integrity of the system.

Failures and Flaws

LRO had very few problems at the system level. The biggest problems during I&T were the CSS interface circuit and solar array gimbal power issues mentioned earlier and high-voltage arcing in our LEND instrument. All three of these issues resulted in hardware changes on the spacecraft. We changed resistors in the C&DH for the CSS's. We added a radiator for the solar array gimbal. And we replaced the LEND instrument with its flight spare before launch. In flight, we have seen an arcing issue with LEND,

but the safing we put in during ground testing has helped prevent any damage. The only other major issue with LRO in flight has been a tight blanket on LOLA. When pointed at cold regions of the moon, the outer layer of the blanket shrinks enough to put some tension between the beamsplitter and the telescope, causing some misalignment of the two. The misalignment causes some of the beams to miss the telescope, but the degradation has been small enough to not impact the LOLA measurement objectives.

Lessons

Motivation—Money is not the right motivator for intellectual challenges. LRO's formula for motivation closely followed Daniel Pink's prescription of autonomy, mastery, and purpose; with outstanding results.

Programmatic Constraints—Programmatic constraints have a profound effect on the development. Tight schedules force decisions and drive the direction of some decisions. A tight budget prior to confirmation ensures that a project makes the hard budget choices. Extra money after CDR, while significant work is occurring in parallel, can save schedule and probably money in the long run by keeping one item from delaying the entire project.

Tight Schedules—Even if the schedule is tight, make sound technical choices. Remember the fortune cookie!

Problems—When a problem pops up, wait 24 hours before reacting. This pause gives the engineers time to understand the situation and develop solutions.

Interface Testing—Interface tests save money in the long run. Test early and often.

Parallel Development—Plan early for parallel development and assembly. You can only add more people to save schedule if the system is architected to use the extra manpower.

Decoupled Delivery—Decouple delivery events so that integration can move forward even if one item is late.

The Team—Systems engineering is all about the team. Success depends on the performance of the entire team. Some people on the team will require more effort, but the extra effort is required to get different perspectives. Watch for overloading, especially in those requiring little effort, and definitely in yourself! Be flexible and optimistic.

Over 1300 people contributed to LRO's success. Each one brought her or his unique talents and perspectives to the team. The end results are the spectacular maps and discoveries of the Lunar Reconnaissance Orbiter.

REFERENCES

- [1] Craig R. Tooley and others, "Lunar Reconnaissance Orbiter Mission and Spacecraft Design," *Space Science Reviews*, Springer Netherlands, ISSN 0038-6308 (Print) 1572-9672 (Online), January 14, 2010, DOI 10.1007/s11214-009-9624-4.
- [2] Benjamin T. Greenhagen and others, "Global Silicate Mineralogy of the Moon from the Diviner Lunar Radiometer," *Science*, Vol. 329, September 17, 2010, p. 1507.
- [3] James W. Head III and others, "Global Distribution of Large Lunar Craters: Implications for Resurfacing and Impactor Populations," *Science*, Vol. 329, September 17, 2010, p. 1504-1507.
- [4] LOLA Image of the Week Archive, <http://lunar.gsfc.nasa.gov/lola/2010may.html>
- [5] NASA Goddard press release No. 10-038, April 26, 2010, http://www.nasa.gov/mission_pages/LRO/news/lro-20100426.html
- [6] LROC web site image of the day, October 28, 2009, <http://lroc.sese.asu.edu/news/index.php?/archives/137-Exploring-the-Apollo-17-Site.html>
- [7] NASA press release No. 09-265, November 13, 2009, http://www.nasa.gov/mission_pages/LCROSS/main/prelim_water_results.html
- [8] Daniel Pink, *Drive: The Surprising Truth about What Motivates Us*, New York: Riverhead Books, 2009.

BIOGRAPHY



Dave Everett has led the design, build, and launch of three spacecraft (FAST, WIRE, and LRO), and he was a key player during the launch of three others (SAMPEX, SWAS, and TRACE). His nineteen years at NASA have included substantial experience in the assembly and testing of spacecraft. Between 1999 and 2005, Mr. Everett focused on early mission planning, including the formulation of Global Precipitation Measurement. Starting in September 2005, Mr. Everett led the technical effort for the Lunar Reconnaissance Orbiter (LRO) as the Mission Systems Engineer. Mr. Everett is currently the chief systems engineer for the Heliophysics and Explorers Program Office at Goddard

Mr. Everett has actively supported NASA outreach activities through over 50 speaking engagements. He has received 29 individual awards and 19 group awards for his efforts at NASA, and he has published 15 papers. He earned a BSEE summa cum laude, at Virginia Tech in 1986 and a MSEE at the University of Maryland in 1989. Before he joined NASA in 1991, Mr. Everett worked at Westinghouse Electric Corporation where he was awarded two patents for his designs of RF circuits.